

Data Warehousing and Data Mining

CPS 116
Introduction to Database Systems

Announcements (November 25)

- ❖ Homework #3 graded
 - Pick them up from Ying during her office hours
- ❖ Homework #4 due today
 - Sample solution available next Tuesday
- ❖ Course project demo period: December 8-13
- ❖ Final exam next Saturday, Dec. 13, 7-10pm
 - Again, open book, open notes
 - Focus on the second half of the course
 - Sample final next Tuesday
 - Sample final solution available Thursday

Data integration

- ❖ Data resides in many distributed, heterogeneous OLTP (On-Line Transaction Processing) sources
 - Sales, inventory, customer, ...
 - NC branch, NY branch, CA branch, ...
- ❖ Need to support OLAP (On-Line Analytical Processing) over an integrated view of the data
- ❖ Possible approaches to integration
 - Eager: integrate in advance and store the integrated data at a central repository called the data warehouse
 - Lazy: integrate on demand; process queries over distributed sources—mediated or federated systems

OLTP versus OLAP

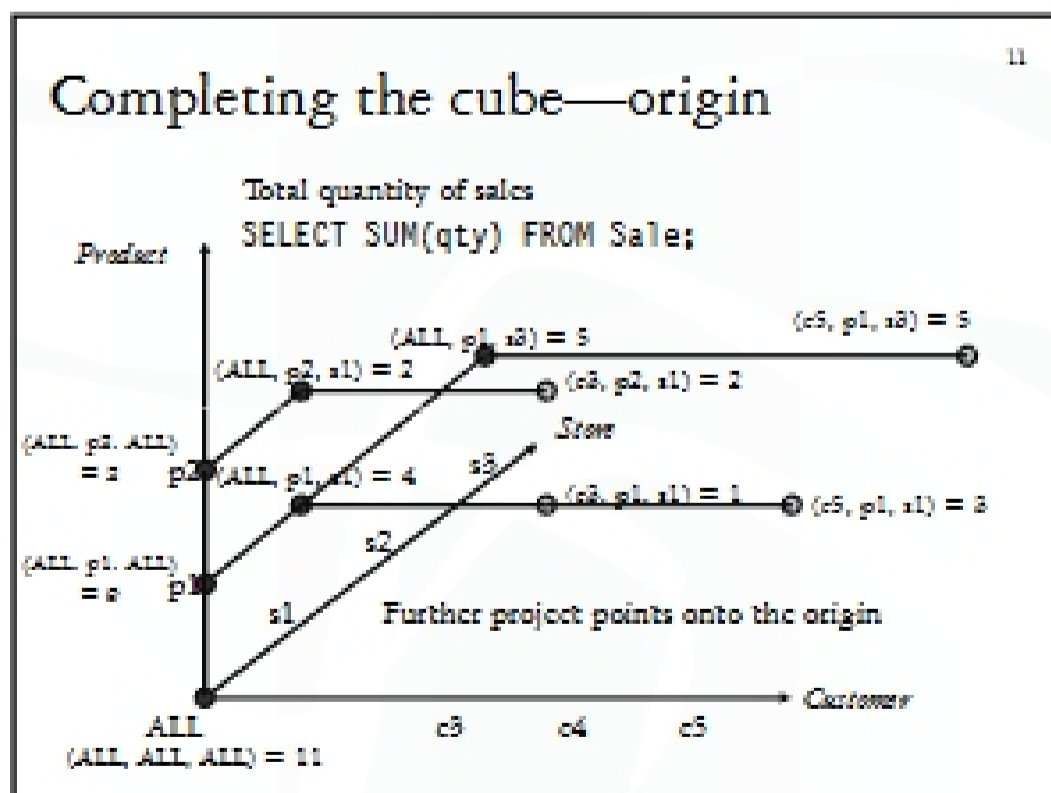
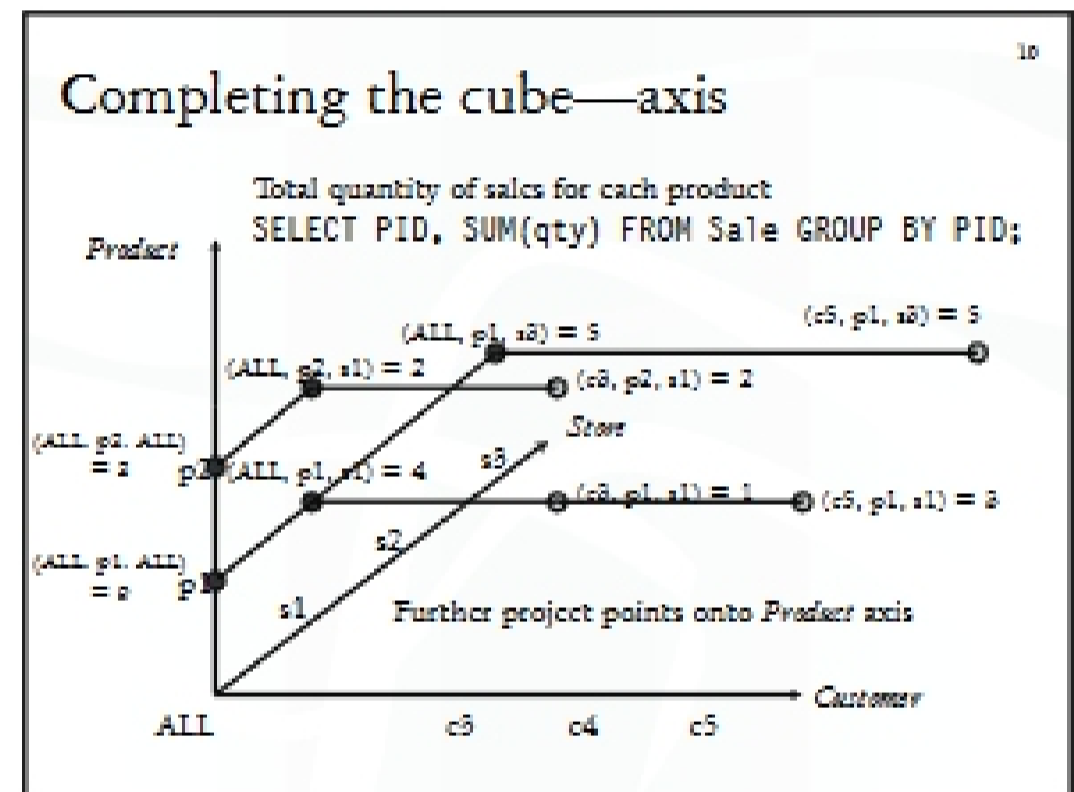
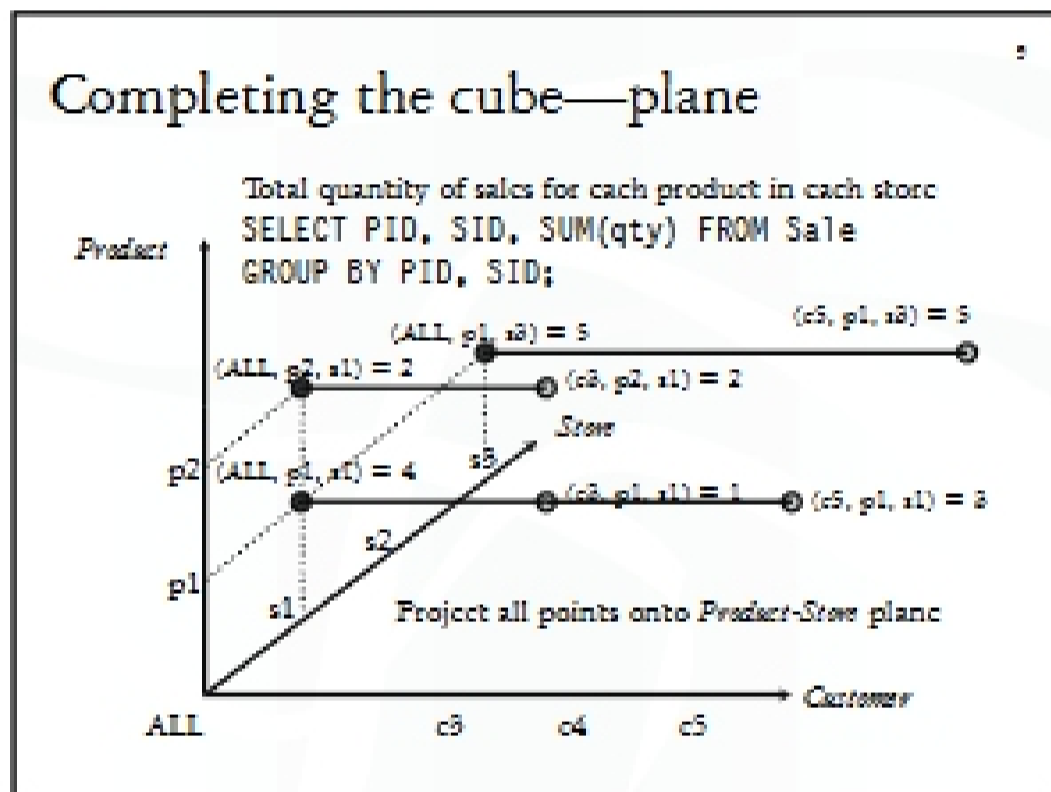
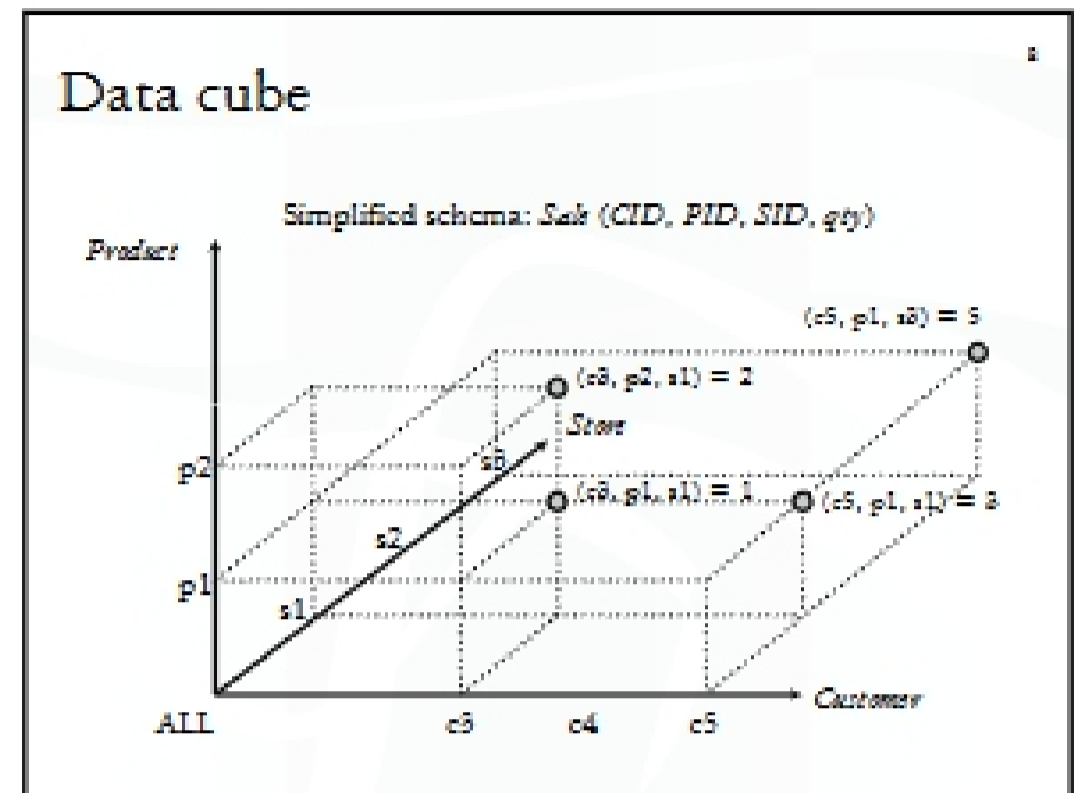
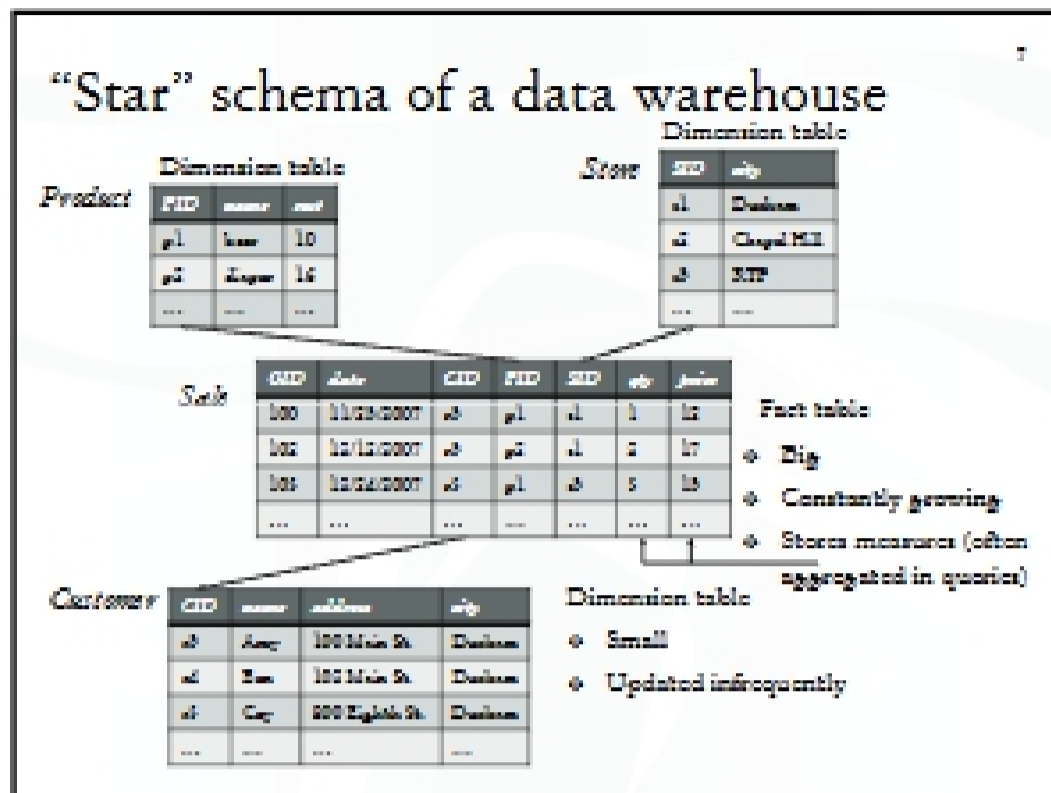
- | OLTP | OLAP |
|--------------------------------------|-----------------------------|
| ❖ Mostly updates | ❖ Mostly reads |
| ❖ Short, simple transactions | ❖ Long, complex queries |
| ❖ Clerical users | ❖ Analysts, decision makers |
| ❖ Goal: ACID, transaction throughput | ❖ Goal: fast queries |
- Implications on database design and optimization?
- OLAP databases do not care much about redundancy
- "Denormalize" tables
 - Many, many indexes
 - Precomputed query results

Eager versus lazy integration

- | Eager (warehousing) | Lazy |
|--|---|
| ❖ In advance: before queries | ❖ On demand: at query time |
| ❖ Copy data from sources | ❖ Leave data at sources |
| ⊖ Answer could be stale | ⊖ Answer is more up-to-date |
| ⊖ Need to maintain consistency | ⊖ No need to maintain consistency |
| ⊖ Query processing is local to the warehouse <ul style="list-style-type: none">▪ Faster▪ Can operate when sources are unavailable | ⊖ Sources participate in query processing <ul style="list-style-type: none">▪ Slower▪ Interferes with local processing |

Maintaining a data warehouse

- ❖ The "ETL" process
 - Extraction: extract relevant data and/or changes from sources
 - Transformation: transform data to match the warehouse schema
 - Loading: integrate data/changes into the warehouse
- ❖ Approaches
 - Recomputation
 - Easy to implement; just take periodic dumps of the sources, say, every night
 - What if there is no "night," e.g., a global organization?
 - What if recomputation takes more than a day?
 - Incremental maintenance
 - Compute and apply only incremental changes
 - Fast if changes are small
 - Not easy to do for complicated transformations
 - Need to detect incremental changes at the sources



CUBE operator

12

- ◆ *Sale* (CID, PID, SID, qty)
- ◆ Proposed SQL extension:


```
SELECT SUM(qty) FROM Sale
GROUP BY CUBE CID, PID, SID;
```
- ◆ Output contains:
 - Normal groups produced by GROUP BY
 - (c1, p1, s1, sum), (c1, p2, s3, sum), etc.
 - Groups with one or more ALL's
 - (ALL, p1, s1, sum), (c2, ALL, ALL, sum), (ALL, ALL, ALL, sum), etc.
- ◆ Can you write a CUBE query using only GROUP BY's?

Gray et al., "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Total." ICDE 1996

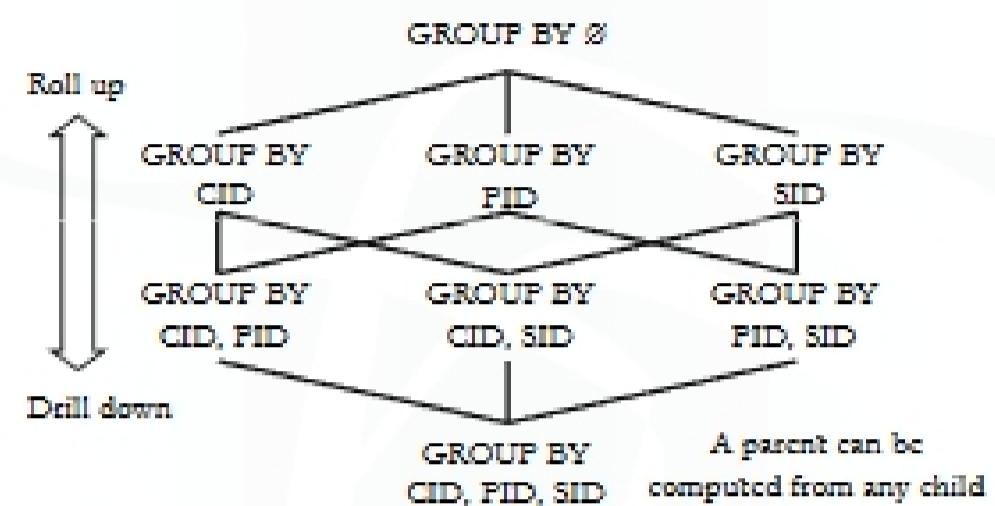
Automatic summary tables

13

- ❖ Computing GROUP BY and CUBE aggregates is expensive
- ❖ OLAP queries perform these operations over and over again
- ❖ Idea: precompute and store the aggregates as automatic summary tables (a DB2 term)
 - Maintained automatically as base data changes
 - Same as materialized views

Aggregation view lattice

14



Selecting views to materialize

15

- ❖ Factors in deciding what to materialize
 - What is its storage cost?
 - What is its update cost?
 - Which queries can benefit from it?
 - How much can a query benefit from it?
- ❖ Example
 - GROUP BY ∅ is small, but not useful to most queries
 - GROUP BY CID, PID, SID is useful to any query, but too large to be beneficial

Harinarayan et al., "Implementing Data Cubes Efficiently." SIGMOD 1996

Data mining

16

- ❖ Data → knowledge
- ❖ DBMS meets AI and statistics
- ❖ Clustering, prediction (classification and regression), association analysis, outlier analysis, evolution analysis, etc.
 - Usually complex statistical "queries" that are difficult to answer → often specialized algorithms outside DBMS
- ❖ We will focus on frequent itemset mining

Mining frequent itemsets

17

- ❖ Given: a large database of transactions, each containing a set of items
 - Example: market baskets
- ❖ Find all frequent itemsets
 - A set of items X is frequent if no less than s_{\min} % of all transactions contain X
 - Examples: {diaper, beer}, {scanner, color printer}

tid	items
T001	diaper milk candy
T002	milk egg
T003	milk beer
T004	diaper milk egg
T005	diaper beer
T006	milk beer
T007	diaper beer
T008	diaper milk beer candy
T009	diaper milk beer
...	...

First try

18

- ❖ A naïve algorithm
 - Keep a running count for each possible itemset
 - For each transaction T , and for each itemset X , if T contains X then increment the count for X
 - Return itemsets with large enough counts
- ❖ Problem: The number of itemsets is huge!
 - 2^x , where x is the number of items
- ❖ Think: How do we prune the search space?